


AD-A183 472 LEARNING INDICES FOR CONCEPTUAL INFORMATION RETRIEVAL: 1/1
AN APPLICATION OF... (U) ILLINOIS UNIV AT URBANA
COORDINATED SCIENCE LAB R MOONEY ET AL. MAY 87
UNCLASSIFIED UILU-ENG-87-2230 N00014-86-K-0309 F/G 5/2 NL

END
9-87
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

COORDINATED SCIENCE LABORATORY
College of Engineering

(2)

AD-A183 472

**LEARNING INDICES FOR
CONCEPTUAL INFORMATION
RETRIEVAL: AN APPLICATION
OF EXPLANATION-BASED
LEARNING IN NATURAL
LANGUAGE PROCESSING**

**Raymond Mooney
Gerald DeJong**

DTIC
ELECTE
JUL 24 1987
S A D

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS None	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) UILLU-ENG-87-2230			5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Coordinated Science Lab University of Illinois		6b. OFFICE SYMBOL (If applicable) N/A	7a. NAME OF MONITORING ORGANIZATION Office of Naval Research	
6c. ADDRESS (City, State, and ZIP Code) 1101 W. Springfield Ave. Urbana, IL 61801			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy St. Arlington, VA 22217	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Office of Naval Research		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014 -86-K-0309	
8c. ADDRESS (City, State, and ZIP Code) 800 N. Quincy St. Arlington, VA 22217			10. SOURCE OF FUNDING NUMBERS	
			PROGRAM ELEMENT NO.	PROJECT NO.
			TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) Learning Indices for Conceptual Information Retrieval: An Application of Explanation-Based Learning in Natural Language Processing				
12. PERSONAL AUTHOR(S) Mooney, Raymond and DeJong, Gerald				
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM TO	14. DATE OF REPORT (Year, Month, Day) May 1987	15. PAGE COUNT 8	
16. SUPPLEMENTARY NOTATION				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Conceptual Information Retrieval, Natural Language Processing, Explanation-Based Learning	
FIELD	GROUP	SUB-GROUP		
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Robust natural language processing systems for conceptual information retrieval require a large number of schemata. For both practical and theoretical reasons, a system cannot be initially programmed with all the schemata it requires. It is therefore important for such a system to be able to learn new schemata automatically during its normal operation. This paper describes the ability of GENESIS, a prototype explanation-based learning system for narrative processing, to use the schemata it learns to index and retrieve specific past episodes. An example run is given which illustrates GENESIS's ability to index and retrieve instances of newly learned schemata.				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL			22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL

Technical Report
UILU-ENG-87-2230

Learning Indices for Conceptual Information Retrieval: An Application of Explanation-Based Learning in Natural Language Processing^{*}

Raymond Mooney
Gerald DeJong

Artificial Intelligence Research Group
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801

February 1987

Accession No.	
NTIS ADAMS	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
1A1	

ABSTRACT

Robust natural language processing systems for conceptual information retrieval require a large number of schemata. For both practical and theoretical reasons, a system cannot be initially programmed with all the schemata it requires. It is therefore important for such a system to be able to learn new schemata automatically during its normal operation. This paper describes the ability of GENESIS, a prototype explanation-based learning system for narrative processing, to use the schemata it learns to index and retrieve specific past episodes. An example run is given which illustrates GENESIS's ability to index and retrieve instances of newly learned schemata.

^{*} This research was supported by the Office of Naval Research under grant N-00014-86-K 0309.

1. Introduction

Conceptual information retrieval involves indexing and retrieving textual information based on an interpretation of its "meaning." As discussed in [Schank81], this approach has a number of advantages over standard information retrieval systems based on keywords. However, conceptual information retrieval requires the ability to "understand" natural language text, which is a very difficult task requiring a large amount of world knowledge. Systems for understanding natural language text (for example [Cullingford78, DeJong82, Dyer83]) generally encode relevant world knowledge in terms of *scripts* or *schemata* [Schank77]. The amount of world knowledge represented in terms of schemata largely determines the performance of such a system. Experience with the FRUMP system [DeJong82] indicated that robustness of a text understanding system is directly related to the number of schemata it possesses.

However, anticipating and encoding all the schemata required for a robust natural language system is impossible for both theoretical and practical reasons. Theoretically, texts can display novel patterns unknown to the implementors of a natural language system. If the natural language system is to respond properly, it must discover such new concepts automatically. Practically, the number of schemata required to cover most natural language domains is prohibitively large and prevents manual programming of all of the necessary concepts. Once again, automatic schema acquisition is essential.

To make the problem more concrete, consider the following example. Suppose we are interested in an automatically updated data base of international trade news stories. We connect the Associated Press newswire to our computer which analyzes trade news stories and files them away for future reference. Conceptually similar stories should be filed together so that a human can easily find relevant information. When a news story arrives, it is analyzed for its conceptual content. This conceptual representation is then added to the data base. Such conceptual processing, while difficult, is becoming increasingly well understood in artificial intelligence, and experimental systems have been constructed for automated conceptual data base updating, for example the CyFr system [Schank81]. The indices for each story are the important schemata that are reported in the story. In most current artificial intelligence systems, the entire range of schemata is defined by a human when the system is implemented. This means that for our conceptual data base system, all possible indices are defined by the time the system sees its first story.

This is unfortunate because input texts can sometimes illustrate novel patterns that themselves should be made into index items. Consider a story about Japan dumping steel in the USA. The story reports that the price of Japanese steel is below Japan's production cost. If the system implementors did not think of the possibility of intentionally selling a product at a loss, the system will flounder on the story. The system may recognize the story as an instance of international trade between the USA and Japan, and that the product is steel which Japan is producing and the USA is purchasing. It might note that the price is unnaturally low. But it will fundamentally misunderstand the point and, therefore, misclassify the story. When a human reader is presented with such a story, he can recognize why dumping steel is profitable in the long run for the Japanese government. It provides fuller employment at home, it pacifies labor, it acquires foreign currency, it stimulates the local economy, and it generates a larger national tax base. In his normal course of processing the story, the human reader has constructed the new concept of "dumping" a product. If the reader later encounters a story about Brazil selling processed sugar beets to Canada at a loss, he can immediately draw a parallel to the Japanese steel story. This is because both stories, for him, are instances of the same important concept, even though they are superficially very different.

It is extremely difficult for current computer systems to index these items together in a data base. A computer system, unable to acquire a concept for dumping a product, cannot recognize the important similarities. For a computer system, the stories look very different: the countries involved are different, the products are different, and the prices (while both abnormally low) are very different. In fact, it will never occur to the computer system that these stories should be compared. There is no more similarity between these stories than between thousands of others.

This is a major flaw of current natural language systems. They cannot acquire new concepts for themselves. The example system cannot augment its set of conceptual indices automatically.

This paper concerns a prototype system called GENESIS which acquires new schemata in the normal course of processing natural language narratives. These schemata are used to improve future processing and also to index future instances so that the system can notice similarities among texts unforeseen by the system's implementors. The GENESIS system is an *explanation-based* learning system [DeJong86, Mitchell86] which means that it acquires a new schema by analyzing the causal structure of a single specific example. Since GENESIS's learning and understanding abilities have been described elsewhere (e.g. [Mooney85a, Mooney85b, Mooney86]), this paper focuses on its recently added ability to use the schemata it learns to index and retrieve specific instances.

2. Similarity-Based Learning Systems in Natural Language Processing

There has already been some work in learning schemata for natural language processing. Both IPP [Lebowitz80] and CYRUS [Kolodner84] learn specializations of existing schemata by analyzing the similarities among a number of examples. For example, IPP started with a general schema for kidnapping and after processing several stories which describe kidnappings in Italy carried out by the terrorist group the Red Brigades, it created a specialized schema for kidnappings in Italy in which the Red Brigades is the default kidnapper. Later, when IPP encounters an article describing a kidnapping in Italy in which the kidnappers are not mentioned, it assumes the Red Brigades is the responsible party. CYRUS also learned specializations in a similarity-based manner and, in addition, used them to index and retrieve specific events. For example, CYRUS started with a general schema for diplomatic meetings and then learned a specialization in which military aid was the topic of discussion. Specific episodes involving meetings about military aid were then indexed under this new schema, and this indexing was used to retrieve answers to questions such as: "Who has Vance talked to about military aid?"

Learning new concepts by detecting and analyzing the similarities and differences among a number of examples is an important and well-researched area in machine learning [Michalski83a, Michalski83b, Quinlan86, Winston75]. However, as discussed in several recent publications (e.g. [Mitchell86, Murphy85]), there are a number of problems and inefficiencies with this approach. The main problems are that it requires a relatively large number of representative examples and fails to take advantage of existing domain knowledge. If the examples encountered by such a system are not representative, the concepts they learn can incorporate spurious correlations which are not a legitimate part of the concept. For example, if all the examples of "dumping of a product" given to a similarity-based system just happen to involve the sale of steel, it will most likely consider steel to be an important part of the new concept.

In addition, both IPP and CYRUS learned specializations of existing schemata and could not learn schemata which were novel *combinations* of existing schemata. In text comprehension, many new concepts involve combining known actions together in a novel way in order to achieve a goal. PAM [Wilensky78] was a plan-based text comprehension system which was capable of understanding such situations; however, it did not learn from its experience in order to improve future performance or to conceptually index and retrieve instances.

3. GENESIS Overview

Unlike CYRUS or IPP, GENESIS is an explanation-based learning system which learns a plan schema from a single instance by determining *why* a particular sequence of actions encountered in a specific story allowed the actors to achieve their goals. The specific instance is then generalized into a schema by removing all of the properties, actions, and relations which do not contribute to this causal explanation. A formal description of the generalization algorithm is given in [Mooney86].

Figure 1 illustrates the overall organization of the system. First, the *parser*, a modified version of McDYPAR [Dyer83], parses English text into assertions in predicate calculus. These

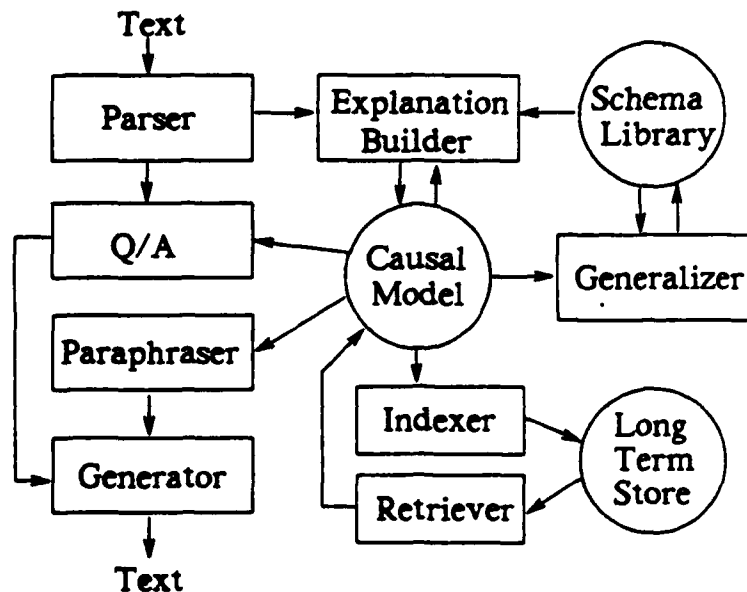


Figure 1: GENESIS System Organization

assertions represent information about the actions, states, and objects in the text. These representations are interpreted by the *explanation-builder* which builds a causal model of the text. The explanation builder attempts to construct explanations for observed actions by causally connecting them to other actions and to characters' goals. This module employs a combination of plan-based [Wilensky78, Wilensky83] and script-based [Cullingford78, DeJong82] understanding mechanisms which access plan schemata stored in the *schema library*. If a character in a narrative achieves an important goal through a novel combination of actions, the *generalizer* generalizes this combination of actions into a new schema. The learned schema is then stored in the schema library where it is available to aid the processing of future texts.

GENESIS also has small modules for answering questions and paraphrasing narratives. The *question-answerer* analyzes explanations in the causal model in order to answer questions about how and why characters performed certain actions. The *paraphraser* uses the most comprehensive schemata detected in the narrative in order to construct a paraphrase. Both of these modules use the *generator* to translate predicate calculus assertions into English text.

After processing a narrative, the *indexer* stores the causal model constructed for this piece of text in the *long-term store* and indexes it under the most comprehensive schemata used in interpreting the story. When answering questions about a particular narrative, the *retriever* can be instructed to retrieve past episodes which are indexed under the same schema used to interpret the current text. These modules allow GENESIS to function as a conceptual retrieval system which, during normal operation, automatically learns new ways to index events.

4. Example Operation of GENESIS

This section presents a sample run of GENESIS which illustrates its ability to learn schemata from a single instance and to use them in indexing and retrieving specific episodes.

First, the system learns schemata for "kidnapping for ransom" and "murder for inheritance." In each case, it learns the schema from a single narrative and saves the causal model it built for the specific episode, indexing it under the new schema. A trace of the system learning these two schemata follows. The names the system gives to the new schemata (MurderInherit and CaptureBargain) are formed simply by concatenating the names of the major actions in the schema.

45_(ProcessStory STORY14)

Input: Claudius owned an estate. Agrippina gave him a poisonous mushroom and he died. Agrippina inherited the estate.

Thematic goal achieved: Agrippina is happy that Agrippina has the estate.

Explanation suitable for generalization. Pruning.. Generalizing... Packaging...

Creating New Schema: (MurderInherit ?x40 ?y10 ?d2)

?x40 is a character. ?d2 is a person. ?x40 murders ?d2. ?y10 is a valuable. ?d2 has ?y10. ?x40 is ?d2's heir. ?x40 inherits ?d2's ?y10.

Having: ?x40 murdered ?d2. suggest MurderInherit

Remembering story as MurderInheritStory1

Ready for questions:

>

46_(ProcessStory STORY1)

Input: Fred is Mary's father and is a millionaire. John approached Mary and pointed a gun at her. She was wearing blue jeans. He told her if she did not get in his car then he would shoot her. He drove her to his hotel and locked her in his room. John called Fred and told him John was holding Mary captive John told Fred if Fred gave him 250000 dollars at Trenos then John would release Mary. Fred gave him the money and John released Mary.

Thematic goal achieved: John is happy that John has the \$250000.

Explanation suitable for generalization. Pruning.. Generalizing...Packaging...

Creating New Schema: (CaptureBargain ?x97 ?a52 ?b11 ?c4 ?y15 ?i19)

?b11 is a person. ?c4 is a location. ?x97 is a character. ?b11 is free. ?x97 captures ?b11. ?a52 is a character. ?x97 contacts ?a52 and tells it that ?b11 is ?x97's captive. ?y15 is a valuable. ?x97 wants to have ?y15 more than it wants ?b11 to be ?x97's captive. ?a52 has a positive relationship with ?b11. ?a52 has ?y15. ?x97 and ?a52 make a bargain in which ?x97 releases ?b11 and ?a52 gives ?x97 ?y15 at ?i19.

Having: ?x97 captured ?b11. suggest CaptureBargain

Having: ?x97 contacted ?a52 and told it that ?b11 was ?x97's captive. suggest CaptureBargain

Having: ?x97 and ?a52 made a bargain in which ?x97 released ?b11 and ?a52 gave ?x97 ?y15 at ?i19. suggest CaptureBargain

Remembering story as CaptureBargainStory1

Ready for questions:

>

After learning these schemata, the system is given two more "murder for inheritance" stories and one more "kidnapping for ransom" story. Each of these narratives is interpreted as an instance of one of the newly learned schemata and its causal model is saved in the long-term store, indexed under that schema. After processing the final instance of each of the new schemata, GENESIS is instructed to: "Review similar stories." This causes the system to retrieve past episodes which are indexed under the same schema used to interpret the present story and make them available for question answering and paraphrasing. Reviewing is done by temporarily replacing the causal model of the current text with the causal model previously constructed and saved for the episode being reviewed.

47_(ProcessStory STORY15)

Input: Gene is Martha's husband and is a millionaire. She shot him and he died. Martha got 1000000 dollars.

Thematic goal achieved: Martha is happy that Martha has the \$1000000.

Remembering story as MurderInheritStory2

Ready for questions:

>

48_(ProcessStory STORY2)

Input: Ted is Alice's husband. He won 100000 dollars in the lottery.

Thematic goal achieved: Ted is happy that Ted has the \$100000.

Bob imprisoned Alice in his basement. Bob got 75000 dollars and released Alice.

Thematic goal achieved: Bob is happy that Bob has the \$75000.

Remembering story as CaptureBargainStory2

Ready for questions:

> Summarize story.

Alice was free. Bob captured Alice. Bob contacted Ted and told him that Alice was Bob's captive. Bob wanted to have \$75000 more than he wanted Alice to be Bob's captive. Ted had a positive relationship with Alice. Ted had the \$75000. Bob and Ted made a bargain in which Bob released Alice and Ted gave Bob the \$75000.

> Where did Bob confine Alice?

Bob confined Alice in the basement.

> Why did Ted want Alice free?

Because Ted was Alice's husband.

> Review similar stories.

There is one other instance of this schema. Review it? yes

Reviewing CaptureBargainStory1

Ready for questions:

> Summarize story.

Mary was free. John captured Mary. John contacted Fred and told him that Mary was John's captive. John wanted to have \$250000 more than he wanted Mary to be John's captive. Fred had a positive relationship with Mary. Fred had the \$250000. John and Fred made a bargain in which John released Mary and Fred gave John the \$250000 at Trenos.

> Why did Fred want Mary free?

Because Fred was Mary's father.

> Where did John confine Mary?

John confined Mary in the room.

> How did John communicate to Fred?

John called Fred and told him that Mary was John's captive.

>

Review finished.

> 49_(ProcessStory STORY17)

Input: Mary is the Bob's mother and is a millionaire. Someone strangled Mary. Bob got \$1000000.

Thematic goal achieved: Bob is happy that Bob has the \$1000000.

Remembering story as MurderInheritStory3

Ready for questions:

> Summarize story.

Bob murdered Mary. Mary had \$1000000. Bob was Mary's heir. Bob inherited Mary's \$1000000.

> How did Bob murder Mary?

Bob strangled Mary.

> Why was Bob Mary's heir?

Because Mary was Bob's mother.

> Review similar stories

There are 2 other instances of this schema. Enter number of story to be reviewed > 1

Reviewing MurderInheritStory1

Ready for questions:

> Summarize story

Agrippina murdered Claudius. Claudius had an estate. Agrippina was Claudius's heir. Agrippina inherited Claudius's estate.

> How did Agrippina murder Claudius?

Agrippina poisoned Claudius with the mushroom.

>

There are 2 other instances of this schema. Enter number of story to be reviewed > 2

Reviewing MurderInheritStory2

Ready for questions:

> Summarize story

Martha murdered Gene. Gene had \$1000000. Martha was Gene's heir. Martha inherited Gene's \$1000000.

> How did Martha murder Gene?

Martha shot Gene and killed him.

> Why was Martha Gene's heir?

Because Gene was Martha's husband.

>

There are 2 other instances of this schema. Enter number of story to be reviewed >

Review finished.

GENESIS is also capable of learning new words which describe learned schemata and using these words to retrieve instances of their associated schemata. GENESIS's word learning abilities are described in more detail in [Mooney87]. Here we will simply present a trace which illustrates the system's ability to use learned words as another index for retrieving past episodes. First, the system is given a "kidnapping for ransom" story in which the word "kidnap" is explicitly mentioned. Given the surrounding context and the use of the word, GENESIS infers that "kidnap" refers to its CaptureBargain schema (see [Mooney87] for details of this process). The system is then given another story about an inheritance¹, and then asked to: "Review kidnapping stories." Given its inferred knowledge about the word "kidnap" and some knowledge of English morphology, it interprets this as a command to review past instances of its CaptureBargain schema.

62_ (ProcessStory STORY8)

Input: Gene is Martha's husband and is Jane's father. Tom locked Jane in his basement. He got 75000 dollars and released Jane.

¹This narrative was originally constructed to present a situation which does not activate the system's learned Murder-Inherit schema even though it refers to a death and an inheritance. Notice that this episode is not stored as an instance of MurderInherit.

Thematic goal achieved: Tom is happy that Tom has the \$75000.

Gene told Martha that someone kidnapped Jane.

Remembering story as CaptureBargainStory3

Ready for questions:

>

63_(ProcessStory STORY16)

Input: Mike is Jan's husband and is a millionaire. Mike died. Jan inherited 1000000 dollars.

Thematic goal achieved: Jan is happy that Jan has the \$1000000.

Ready for questions:

> Summarize

Mike died. Jan inherited Mike's \$1000000.

> Review kidnapping stories.

There are 3 instances of CaptureBargain. Enter number of story to be reviewed > 3

Reviewing CaptureBargainStory3

Ready for questions:

> Summarize

Jane was free. Tom captured Jane. Tom contacted Gene and told him that Jane was Tom's captive. Tom wanted to have \$75000 more than he wanted Jane to be Tom's captive. Gene had a positive relationship with Jane. Gene had the \$75000. Tom and Gene made a bargain in which Tom released Jane and Gene gave Tom the \$75000.

>

There are 3 instances of CaptureBargain. Enter number of story to be reviewed >

Review finished.

5. Conclusions

We have argued that it is important for a natural language processing system for conceptual information retrieval to be able to learn new schemata during its normal course of its operation. Previous text processing systems which learned schemata, such as IPP and CYRUS, used similarity-based learning techniques and could learn only specializations of existing schemata. GENESIS is an explanation-based learning system which learns schemata defined by a novel combination of existing plans. It uses its learned schemata to index and retrieve specific past episodes as well as to improve its ability to process narratives. These features make it interesting as an initial prototype for a robust conceptual information retrieval system which improves its performance with experience and is able to accommodate unforeseen situations.

References

- [Cullingford78] R. E. Cullingford, "Script Application: Computer Understanding of Newspaper Stories," Technical Report 116, Department of Computer Science, Yale University, New Haven, CT, January 1978.
- [DeJong82] G. F. DeJong, "An Overview of the FRUMP System," in *Strategies for Natural Language Processing*, W. G. Lehnert and M. H. Ringle (ed.), Lawrence Erlbaum and Associates, Hillsdale, NJ, 1982.
- [DeJong86] G. F. DeJong and R. J. Mooney, "Explanation-Based Learning: An Alternative View," *Machine Learning* 1, 2 (1986), pp. 145-176.
- [Dyer83] M. J. Dyer, *In-Depth Understanding*, MIT Press, Cambridge, MA, 1983.
- [Kolodner84] J. L. Kolodner, *Retrieval and Organization Strategies in Conceptual Memory*, Lawrence Erlbaum and Associates, Hillsdale, NJ, 1984.
- [Lebowitz80] M. Lebowitz, "Generalization and Memory in an Integrated Understanding System," Technical Report 186, Ph.D Thesis, Department of Computer Science, Yale University, New Haven, CT, 1980.
- [Michalski83a] R. S. Michalski and R. E. Stepp, "Learning from Observation: Conceptual Clustering," in *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell and T. M. Mitchell (ed.), Tioga Publishing Company, Palo Alto, CA, 1983, pp. 331-363.
- [Michalski83b] R. S. Michalski, "A Theory and Methodology of Inductive Learning," in *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (ed.), Tioga Publishing Company, Palo Alto, CA, 1983, pp. 83-134.
- [Mitchell86] T. M. Mitchell, R. Keller and S. Kedar-Cabelli, "Explanation-Based Generalization: A Unifying View," *Machine Learning* 1, 1 (January 1986), pp. 47-80.
- [Mooney85a] R. J. Mooney and G. F. DeJong, "Learning Schemata for Natural Language Processing," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985, pp. 681-687.
- [Mooney85b] R. J. Mooney, "Generalizing Explanations of Narratives into Schemata," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, May 1985. (Also appears as Technical Report T-159, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.)
- [Mooney86] R. J. Mooney and S. W. Bennett, "A Domain Independent Explanation-Based Generalizer," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, August 1986, pp. 551-555. (A longer updated version appears as Technical Report UILU-ENG-86-2216, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign)
- [Mooney87] R. J. Mooney, "Integrated Learning of Words and their Underlying Concepts," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, July 1987. (Also appears as Technical Report UILU-ENG-87-2229, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign)
- [Murphy85] G. L. Murphy and D. L. Medin, "The Role of Theories in Conceptual Coherence," *Psychological Review* 92, 3 (July 1985), pp. 289-316.
- [Quinlan86] J. R. Quinlan, "Induction of Decision Trees," *Machine Learning* 1, 1 (1986), pp. 81-106.
- [Schank77] R. C. Schank and R. P. Abelson, *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*, Lawrence Erlbaum and Associates, Hillsdale, NJ, 1977.

- [Schank81] R. C. Schank, J. L. Kolodner and G. F. DeJong, "Conceptual Information Retrieval," in *Information Retrieval Research*, R. N. Oddy, S. E. Robertson, C. J. van Rijsbergen, and P. W. Williams (ed.), Butterworths, London, 1981.
- [Wilensky78] R. W. Wilensky, "Understanding Goal-Based Stories," Technical Report 140, Ph.D. Thesis, Department of Computer Science, Yale University, New Haven, CT, September 1978.
- [Wilensky83] R. W. Wilensky, *Planning and Understanding: A Computational Approach to Human Reasoning*, Addison-Wesley, Reading, MA, 1983.
- [Winston75] P. H. Winston, "Learning Structural Descriptions from Examples," in *The Psychology of Computer Vision*, P. H. Winston (ed.), McGraw-Hill, New York, NY, 1975, pp. 157-210.

END

9-87

DTIC